

Deformable Organisms for Automatic Medical Image Analysis

Ghassan Hamarneh^{1,2}, Tim McInerney^{3,2}, and Demetri Terzopoulos^{4,2}

¹ Dept. of Signals and Systems, Chalmers University of Technology, Göteborg 41296, Sweden
ghassan@s2.chalmers.se, www.s2.chalmers.se

² Dept. of Computer Science, University of Toronto, Toronto M5S 3H5, Canada
{ghassan,tim,dt}@cs.toronto.edu, www.cs.toronto.edu

³ School of Computer Science, Ryerson University, Toronto M5B 2K3, Canada
tmcinern@scs.ryerson.ca, www.scs.ryerson.ca

⁴ Courant Institute, New York University, New York NY 10003, USA
dt@cs.nyu.edu, www.mrl.nyu.edu

Abstract. We introduce a new paradigm for automatic medical image analysis that adopts concepts from the field of Artificial Life. Our approach prescribes deformable organisms, autonomous agents whose objective is the segmentation and analysis of anatomical structures in medical images. A deformable organism is structured as a ‘muscle’-actuated ‘body’ whose behavior is controlled by a ‘brain’ that is capable of making both reactive and deliberate decisions. This intelligent deformable model possesses an ‘awareness’ of the segmentation process, which emerges from a conflux of perceived sensory data, an internal mental state, memorized knowledge, and a cognitive plan. We develop a class of deformable organisms using a medial representation of body morphology that facilitates a variety of controlled local deformations at multiple spatial scales. Specifically, we demonstrate a deformable ‘worm’ organism that can overcome noise, incomplete edges, considerable anatomical variation, and occlusion in order to segment and label the corpus callosum in 2D mid-sagittal MR images of the brain.

1 Introduction

The automatic segmentation and labeling of anatomical structures in medical images is a persistent problem that continues to defy solution. There is consensus within the medical image analysis research community that the development of general-purpose automatic segmentation algorithms will require not only powerful bottom-up, data-driven processes, but also equally powerful top-down, knowledge-driven processes within a robust decision-making framework that operates across multiple levels of abstraction [2]. Deformable models, one of the most actively researched model-based segmentation techniques [5], feature a potent bottom-up component founded in estimation theory, optimization, and physics-based dynamical systems, but their top-down processes have traditionally relied on interactive initialization and guidance by knowledgeable users. Attempts to fully automate deformable model segmentation methods have so far been less than successful at coping with the enormous variation in anatomical structures of interest, the significant variability of image data, the need for intelligent initialization conditions, etc.

The time has come to shift our attention to what promises to be a critical element in any viable, highly automated solution: the decision-making framework itself. Existing decision-making strategies for deformable models are inflexible and do not operate at an appropriate level of abstraction. Hierarchically organized models, which shift their focus from structures associated with stable image features to those associated with less stable features, are a step in the right direction [4,9]. However, high-level contextual

knowledge remains largely ineffective because it is intertwined much too tightly with the low-level optimization-based mechanisms. It is difficult to obtain intelligent, global (i.e., over the whole image) model behavior throughout the segmentation process from such mechanisms. In essence, current deformable models have no explicit awareness of where they (or their parts) are in the image or what their objectives are at any time during the optimization process.

It is our contention that we must revisit ideas for incorporating knowledge that were explored in earlier systems (e.g., [14]), and develop new algorithms that focus on top-down reasoning strategies which may best leverage the powerful bottom-up feature detection and integration abilities of deformable models and other modern model-based medical image analysis techniques. We further contend that a layered architecture is appropriate, where the high-level reasoning layer has knowledge about and control over the low-level model (or models) at all times. The reasoning layer should apply an active, explicit search strategy that first looks for the most stable image features before proceeding to less stable image features, and so on. It should utilize contextual knowledge to resolve regions where there is a deficiency of image feature information.

To achieve these goals, we introduce a new paradigm for automatic medical image analysis that adopts concepts from the emerging field of Artificial Life. In particular, we develop *deformable organisms*, autonomous agents whose objective is the segmentation and analysis of anatomical structures in medical images. A deformable organism is structured as a ‘muscle’-actuated ‘body’ whose behavior is controlled by a ‘brain’ that is capable of making both reactive and deliberate decisions. This intelligent deformable model possesses a non-trivial ‘awareness’ of the segmentation process, which emerges from a conflux of perceived sensory data, an internal mental state, memorized knowledge, and a cognitive plan. By constructing deformable organisms in a layered fashion, we are able to separate the knowledge-driven model-fitting control functionality from the data-driven, local image feature integration functionality, exploiting both for maximal effectiveness.

1.1 Artificial Life Modeling

The Artificial Life (ALife) modeling approach has been applied successfully to produce realistic computer graphics models of plants and animals [13]. Artificial animals are relevant to deformable organisms. Autonomous agents known as “artificial fishes” [12] serve to illustrate the key functional components of artificial animals: bodies that comprise muscle actuators, sensory organs (eyes, etc.) and, most importantly, brains consisting of motor, perception, behavior, learning and cognition centers. Controllers in the motor center coordinate muscle actions to carry out specific motor functions, such as locomotion and sensor actuation. The perception center employs attention mechanisms to interpret sensory information about the dynamic environment. The behavior center realizes an adaptive sensorimotor system through a repertoire of behavior routines that couple perception to action in meaningful ways. The learning center in the brain enables the artificial animal to learn motor control and behavior through practice and sensory reinforcement. The cognition center enables it to think.

To manage their complexity, artificial animal models are best organized hierarchically, such that each successive modeling layer augments the more primitive functionalities of lower layers. At the base of the modeling hierarchy (see Fig 1a), a geometric modeling layer represents the morphology and appearance of the animal. Next, a physical modeling layer incorporates biomechanical principles to constrain the




Fig. 1. (a) The ALife modeling pyramid (adapted from [12]). (b) A deformable organism: The brain issues ‘muscle’ actuation and perceptual attention commands. The organism deforms and senses image features, whose characteristics are conveyed to its brain. The brain makes decisions based on sensory input, memorized information and prior knowledge, and a pre-stored plan, which may involve interaction with other organisms.

geometry and simulate biological tissues. Further up the hierarchy is a motor control layer that motivates internal muscle actuators in order to synthesize lifelike locomotion. Behavioral and perceptual modeling layers cooperate to support a reactive behavioral repertoire. At the apex of the modeling pyramid is a cognitive modeling layer, which simulates the deliberative behavior of higher animals, governs what an animal knows about itself and its world, how that knowledge is acquired and represented, and how automated reasoning and planning processes can exploit knowledge to achieve high-level goals.

1.2 An Artificial Life Modeling Paradigm for Medical Image Analysis

Viewed in the context of the artificial life modeling hierarchy (Fig. 1a), current *automatic* deformable model-based approaches to medical image analysis include geometric and physical modeling layers only (in interactive deformable models, such as snakes, the human operator is relied upon to provide suitable behavioral level and cognitive level support). At the physical level, deformable models interpret image data by simulating dynamics or minimizing energy terms, but the models themselves do not monitor or control this optimization process except in a most primitive way. At the geometric level, aside from a few notable exceptions [11], deformable models are not generally designed with intuitive, multi-scale, multi-location deformation ‘handles’. Their inability to perform global deformations, such as bending, and other global motions such as sliding and backing up makes it difficult to develop reasoning or planning strategies for these models at the correct level of abstraction [5].

In more sophisticated deformable models, prior information is used to constrain shape and appearance, as well as the statistical variation of these quantities [1,10]; however, these models have no explicit awareness of where they are and, consequently, the effectiveness of these constraints is dependent upon model starting conditions. The lack of awareness also prevents the models from knowing when to trust the image feature information and ignore the constraint information and vice versa. The constraint information is therefore applied arbitrarily. Furthermore, because there is no active, explicit search for stable image features, the models are prone to latching onto incorrect features [1] simply due to their proximity and local decision-making. Once this latching

occurs, the lack of control of the fitting procedure prevents the model from correcting the misstep. The result is that the local decisions that are made do not add up to intelligent global behavior.

To overcome the aforementioned deficiencies while retaining the core strengths of the deformable model approach, we add high-level controller layers (a ‘brain’) on top of the geometric and physical (or deformation) layers to produce an autonomous deformable organism (Fig. 1b). The intelligent activation of these lower layers allows the organism to control the fitting/optimization procedure. The layered architecture approach allows the deformable organism to make deformation decisions at the correct level of abstraction.

The perception system of the deformable organism comprises a set of sensors that provide information. Any type of sensors can be incorporated, from edge strength and edge direction detectors to snake ‘feelers’. Sensors can be focused or trained for specific image features and image feature variation in a task-specific way; hence, the organism can disregard sensory information superfluous to its current behavioral needs.

Explicit feature search requires powerful, flexible and intuitive model deformation control. We achieve this with a set of ‘motor’ (i.e. deformation) controllers, which are parameterized procedures dedicated to carrying out a complex deformation function, such as successively bending a portion of the organism over some range of angles or stretching part of the organism forward some distance.

The organism is ‘self-aware’ (i.e. knows where it and its parts are and what it is seeking) and therefore it effectively utilizes global contextual knowledge. The organism begins by searching for the most stable anatomical features in the image before proceeding to less stable features. Once stable features are found and labeled, the organism uses neighboring information and prior knowledge to determine the object boundary in regions known to provide little or no feature information.

Because the organism carries out active, explicit searches for object features, it is not satisfied with the nearest matching feature but looks further within a region to find the best match, thus avoiding local minimum solutions. Furthermore, by carrying out explicit searches for features we ensure correct correspondence between the model and the data. If a feature cannot be found, the organism flags the situation. Subsequently, if multiple plans exist, another plan could potentially be selected and the search for the missing feature postponed until further information is available.

2 A Deformable Organism for 2D MR Brain Image Analysis

To demonstrate the potential of our framework for medical image analysis, we have developed a deformable “worm” organism that can overcome noise, incomplete edges, considerable anatomical variation, and occlusion in order to segment and label the corpus callosum (CC) in 2D mid-sagittal MR images of the brain. We will now describe in detail the layered architecture for this particular deformable organism.

2.1 Geometric Representation

As its name suggests, the deformable worm organism is based on a medial representation of body morphology [3] that facilitates a variety of controlled local deformations at multiple spatial scales. In this shape representation scheme, the CC anatomical structure is described with four shape profiles derived from the primary medial axis of the CC boundary contour. The medial profiles describe the geometry of

the structure in a natural way and provide general, intuitive, and independent shape measures. These profiles are: a length profile $L(m)$, an orientation profile $O(m)$, a left (with respect to the medial axis) thickness profile $T^l(m)$, and a right thickness profile $T^r(m)$ where $m = 1, 2, \dots, N$ and N is the number of medial nodes. The length profile represents the distances between consecutive pairs of medial nodes, and the orientation profile represents the angles of the edges connecting the pairs of nodes. The thickness profiles represent the distances between medial nodes and their corresponding boundary points (Fig. 2, Fig. 3)¹.




Fig. 2. (a) CC anatomical feature labels overlaying a reconstruction of the CC using the medial shape profiles shown in Fig. 3. (b) Diagram of shape representation.




Fig. 3. Example medial shape profiles: (a) length, (b) orientation, (c) left and (d) right thickness profiles.

2.2 Motor System

Shape Deformation Actuators. In addition to affine transformation abilities (translate, rotate, scale), we control organism deformation by defining deformation actuators in terms of the medial shape profiles (Fig. 4). Controlled stretch (or compress), bend, and bulge (or squash) deformations are implemented as deformation operators acting on the length, orientation, or thickness profiles, respectively. Furthermore, by utilizing a hierarchical (multiscale) and regional principal component analysis to capture the shape variation statistics in a training set [3], we can keep the deformations consistent with prior knowledge of possible shape variations. Whereas general, statistically-derived shape models produce global shape variation modes only [1,10], we are able to produce spatially-localized feasible deformations at desired scales, thus supporting our goal of intelligent deformation planning.

Several operators of varying types, amplitudes, scales, and locations can be applied to any of the length, orientation, and thickness shape profiles (Fig. 5a-d). Similarly, multiple statistical shape variation modes can be activated, with each mode acting at a specified amplitude, location and scale of the shape profiles (Fig. 5e-h). In general, operator- and statistics-based deformations can be combined (Fig. 5i) and expressed as

$$p_d = \bar{p}_d + \sum_t \sum_s \left(M_{dls} w_{dls} + \sum_t \alpha_{dlst} k_{dlst} \right) \quad (1)$$

¹ Currently we construct medial profiles only from the primary medial axis and have not considered secondary axes. This may prevent the CC worm organism from accurately representing highly asymmetrical (with respect to the primary axis) parts of some corpora callosa. We also realize that our medial shape representation needs improvement near the end caps. We are currently exploring these issues, as well as issues related to the extension of our model to 3D, and we intend to make full use of the considerable body of work of Pizer *et al* [6,7,8] on these topics.

where p is a shape profile, d is a deformation type (stretch, bend, left/right bulge), i.e. $p_d(m) : \{L(m), O(m), T^l(m), T^r(m)\}$, \bar{p} is the average shape profile, k is an operator profile (with unity amplitude), l and s are the location and scale of the deformation, t is the operator type (e.g. Gaussian, triangular, flat, bell, or cusp), α is the operator amplitude, the columns of M are the variation modes for a specific d , l , and s , and w contains variation mode weights. Details can be found in [3].




Fig. 4. Introducing a bulge on the upper boundary of the CC by applying a deformation operator on the upper thickness profile, $T^r(m)$. (a) $T^r(m)$ before and (c) after applying the operator. (b) The reconstructed shape before and (d) after the operator.




Fig. 5. Examples of controlled deformations: (a)-(c) Operator-based bulge deformation at varying locations/amplitudes/scales. (d) Operator-based stretching with varying amplitudes over entire CC. (e)-(g) Statistics-based bending of left end, right end, and left half of CC. (h) Statistics-based bulge of the left and right thickness over entire CC. (i) From left to right: (1) mean shape, (2) statistics-based bending of left half, followed by (3) locally increasing lower thickness using operator, followed by (4) applying operator-based stretch and (5) adding operator based bend to right side of CC.

Deformation (Motor) Controllers. The organism's low-level motor actuators are controlled by motor controllers. These parameterized procedures carry out complex deformation functions such as sweeping over a range of rigid transformation parameters, sweeping over a range of stretch/bend/thickness amplitudes at a certain location and scale, bending at increasing scales, moving a bulge on the boundary etc. Other high-level deformation capabilities include, for example, smoothing the medial/left/right boundaries, interpolating a missing part of the thickness profile, moving the medial axis to a position midway between the left and right boundaries, and re-sampling the model by including more medial and boundary nodes.

2.3 Perception System

Different parts of the organism are dynamically assigned sensing capabilities and thus act as sensory organs (SOs) or receptors. The locations of the SOs are typically confined to the organism's body (on-board SOs) such as at its medial or boundary nodes, at curves or segments connecting different nodes. In our implementation, the SOs are made sensitive to different stimuli such as image intensity, image gradient magnitude and direction, a non-linearly diffused version of the image, an edge detected (using Canny's edge detector) image, or even the result of a Hough transform. In general, a wide variety of image processing/analysis techniques can be applied to the original image.

2.4 Behavioral/Cognitive System

The organism's cognitive center combines sensory information, memorized information, and instructions from a pre-stored segmentation plan to carry out active, explicit searches for object features by activating 'behavior' routines. Behavior routines are designed based on available organism motor skills, perception capabilities, and available anatomical landmarks. For example, the routines implemented for the CC worm organism include: find-top-of-head, find-upper-boundary-of-CC, find-genu, find-rostrum, find-splenium, latch-to-upper-boundary, latch-to-lower-boundary, find-fornix, thicken-right-side, thicken-left-side, back-up. The behavior routines subsequently activate the deformation controllers to complete a stage in the plan and bring the organism closer to its intention of object segmentation.

The segmentation plan provides a means for human experts to incorporate global contextual knowledge. It contains instructions on how best to achieve a correct segmentation by optimally prioritizing behaviors. If we know, for example, that the corner-shaped rostrum of the CC is always very clearly defined in an MRI image, then the find-rostrum behavior should be given a very high priority. Adhering to the segmentation plan and defining it at a behavioral level affords the organism with an awareness of the segmentation process. This enables it to make effective use of prior shape knowledge – it is applied only in anatomical regions of the target object where there is a high level of noise or known gaps in the object boundary edges, etc. In the next section we describe the segmentation plan for the CC organism to illustrate this ability to harness global contextual knowledge.

3 Results

When a CC deformable worm organism is released into a 2D sagittal MRI brain image, it engages in different 'behaviors' as it progresses towards its goal. Since the upper boundary (Fig. 2a) of the CC is very well defined and can be easily located with respect to the top of the head, the cognitive center of the CC organism activates behaviors to first locate the top of the head and then move downwards (through the gray and white matter) in the image space to locate the upper boundary (Fig. 6.1-5). Next, the organism bends to latch to the upper boundary and activates a find-genu routine, causing the CC organism to stretch and grow along this boundary towards the genu (Fig. 6.6-7). Once the genu is located, the find-splenium routine is activated and the organism stretches and grows in the opposite direction (Fig. 6.11). The genu and splenium are easily detected by looking for a sudden change in direction of the upper boundary towards the middle of the head.

Once the genu is found, the organism knows that the lower boundary opposite to the genu is well defined so it backs up and latches to the lower boundary (Fig. 6.8). It then activates the find-rostrum behavior that tracks the lower boundary until it reaches the distinctive rostrum (Fig. 6.8-10). At the splenium end of the CC, the organism backs up and finds the center of a circle that approximates the splenium end cap (Fig. 6.12). The lower boundary is then progressively tracked from the rostrum to the splenium while maintaining parallelism with the organism's medial axis in order to avoid latching to the potentially occluding fornix (Fig. 6.13-14). However, the lower boundary may still dip towards the fornix, so a successive step is performed to locate where, if at all, the fornix occludes the CC, by activating the find-fornix routine (making use of edge strength along the lower boundary, its parallelism to the medial axis, and statistical thickness values). Thus, prior knowledge is applied only when and where required. If the fornix does indeed occlude the CC, any detected dip in the organism's boundary is repaired by interpolating neighboring thickness values. The thickness of the upper boundary is then adjusted to latch on to the corresponding boundary in the image (Fig. 6.15-17). At this point the CC organism has almost reached its goal; however, the medial axis is not in the middle of the CC organism (Fig. 6.18), hence the medial axis is re-parameterized by positioning the medial nodes halfway between the boundary nodes (Fig. 6.19-20). Finally the lower and upper boundaries are re-located again to obtain the final segmentation result (Fig. 6.21).

In addition, Fig. 7 demonstrates the detection and repairing of the fornix. Fig. 8 demonstrates the organism's self-awareness. Fig. 9 shows other segmentation results and several validated examples are also shown in Fig. 10.




Fig. 6. Intelligent CC organism progressing through a sequence of behaviors to segment the CC.




Fig. 7. (a) Before and (b) after detecting and repairing the fornix dip. (c) The gradient magnitude.




Fig. 8. The CC organism’s self-awareness makes it capable of identifying landmark parts.




Fig. 9. Example segmentation results.




Fig. 10. Example segmentation results (top), also shown (in black) over manually segmented (gray) CC (bottom).

4 Conclusions

Robust, automatic medical image analysis requires the incorporation and intelligent utilization of global contextual knowledge. We have introduced a new paradigm for medical image analysis that applies concepts from artificial life modeling to meet this requirement. By architecting a deformable model-based framework in a layered fashion, we are able to separate the ‘global’ model-fitting control functionality from the local feature integration functionality. This separation allows us to define a model-fitting controller or ‘brain’ in terms of the high-level anatomical features of an object rather than low-level image features. The layered-architecture approach also provides the brain layer with precise control over the lower-level model deformation layer. The result is an intelligent organism that is continuously aware of the progress of the segmentation, allowing it to effectively apply prior knowledge of the target object. We have demonstrated the potential of this approach by constructing a Corpus Callosum “worm” organism and releasing it into MRI brain images in order to segment and label the CC.

Several interesting aspects of our approach are currently in consideration for further exploration. These include extending our model to 3D, designing a motion tracking plan and releasing an organism into time-varying image ‘environments’ (i.e. 4D images), exploring the use of multiple plans and plan selection schemes, and exploring the application of learning algorithms, such as genetic algorithms, to assist human experts in the generation of optimal plans. Another potentially important research direction is the use of multiple organisms that intercommunicate contextual image information (i.e. are ‘aware’ of one another).

Acknowledgements

GH was funded in part by the Visual Information Technology (VISIT) program, Swedish Foundation for Strategic Research (SSF). Dr. Martha Shenton of the Harvard Medical School generously provided the MRI data.

References

1. Cootes, T., Beeston, C., Edwards, G., Taylor, C.: A Unified Framework for Atlas Matching using Active Appearance Models. *Image Processing in Medical. Imaging* (1999) 322-333.
2. Duncan, J., Ayache, N.: Medical Image Analysis: Progress Over Two Decades and the Challenges Ahead. *IEEE Trans. PAMI* 22(1) (2000) 85 -106.
3. Hamarneh, G., McInerney, T.: Controlled Shape Deformations via Medial Profiles. *Vision Interface* (2001) 252-258.
4. McInerney, T., Kikinis, R.: An Object-based Volumetric Deformable Atlas for the Improved Localization of Neuroanatomy in MR Images. *MICCAI* (1998) 861-869.
5. McInerney, T., Terzopoulos, D.: Deformable Models in Medical Image Analysis: A Survey. *Medical Image Analysis* 1(2) (1996) 91-108.
6. Pizer, S., Fletcher, P., Fridman, Y., Fritsch D., Gash, A., Glotzer, J., Joshi, S., Thall A., Tracton, G., Yushkevich, P., Chaney, E.: Deformable M-Reps for 3D Medical Image Segmentation. Submitted to *Medical Image Analysis* (2000).
7. Pizer, S., Fritsch, D., Low, K., Furst, J.: 2D & 3D Figural Models of Anatomic Objects from Medical Images. *Mathematical Morphology and Its Applications to Image Processing, Kluwer Computational Imaging and Vision Series, H.J.A.M. Heijmans, J.B.T.M. Roerdink, Eds. Amsterdam, (1998) 139-150.*
8. Pizer, S., Fritsch, D.: Segmentation, Registration, and Measurement of Shape Variation via Image Object Shape. *IEEE Transactions on Medical Imaging* 18(10) (1999) 851-865.
9. Shen, D., Davatzikos, C.: An Adaptive-Focus Deformable Model Using Statistical and Geometric Information. *IEEE Trans. PAMI* 22(8) (2000) 906 -913.
10. Szekely, G., Kelemen, A., Brechbuehler C., Gerig, G.: Segmentation of 3D Objects From MRI Volume Data Using Constrained Elastic Deformations of Flexible Fourier Surface Models. *Medical Image Analysis* 1(1) (1996) 19-34.
11. Terzopoulos, D., Metaxas, D.: Dynamic 3D Models with Local and Global Deformations: Deformable Superquadrics. *IEEE Trans. PAMI* 13(7) (1991) 703-714.
12. Terzopoulos, D., Tu, X., Grzeszczuk, R.: Artificial Fishes: Autonomous Locomotion, Perception, Behavior, and Learning in a Simulated Physical World. *Artificial Life* 1(4) (1994)
13. Terzopoulos, D.: Artificial Life for Computer Graphics. *Commun. ACM* 42(8) (1999) 32-42.
14. Tsotsos, J.K., Mylopoulos, J., Covvey, H.D., Zucker, S.W.: A Framework for Visual Motion Understanding. *IEEE Trans. PAMI* 2(6) (1980) 563-573.